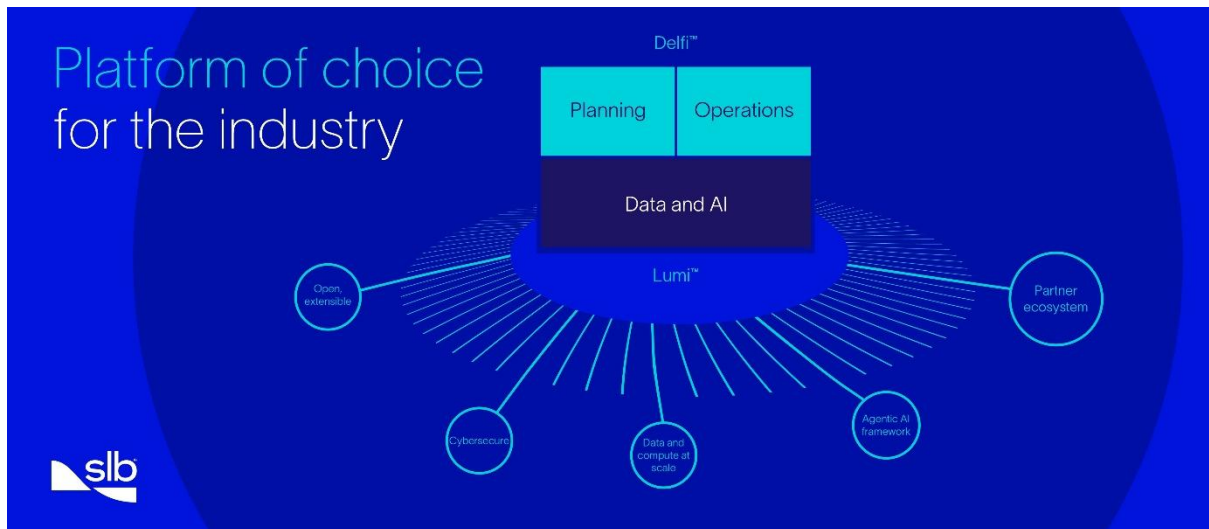


Trygve Randen

Senior Vice President of
Digital Products and Solutions

Thank you Rakesh.



SLB's platform approach has been the cornerstone of our technology strategy for more than two decades. To start, it is worth grounding us in what we mean by platforms because the term is used loosely.

A platform must do two things: make an enterprise's trusted data accessible and provide an open environment in which that data is consumed by applications, by workflows, and increasingly, by AI.

It is the layer that makes everything work together at scale. In our industry, that bar is unusually high.

As Rakesh explained, the data is complex, domain-specific, and often business- and safety-critical. The workflows span multiple scientific and engineering disciplines. The operating environments are global and frequently constrained by data residency, and increasingly, by technology sovereignty. A horizontal, multi-purpose cloud platform does not meet these needs.

What is required are platforms built for the domain, that understand the data, understand the workflows, and can operate at enterprise scale for the most demanding customers.

Our customers operate in a world of multiple vendors and proprietary data. We accommodate customers who wish to bring their intellectual property and run it alongside ours in a governed environment. We even partner with many of our customers to co-develop technology through open APIs and an extensible application framework, a philosophy we pioneered twenty years ago with the industry's first commercial plug-in environment.

Customers and third-party developers can deploy their own applications, algorithms, and proprietary workflows alongside ours in Delfi, Lumi and Tela. They bring their workflows, we provide the platform.

Turning to security: our platforms operate under the tightest standards, with data encrypted in transit and at rest, multi-factor authentication, and role-based access controls.

For an industry that deals in competitively sensitive subsurface data and operates under regulatory oversight, this is not a feature but a prerequisite. There is no doubt we have the data, but one of the structural constraints that has held this industry back is the physical limitations of traditional computing infrastructure.

A reservoir simulation that takes three weeks on a local workstation can run overnight using elastic cloud resources. A seismic processing job that would require a dedicated data center can be scaled on demand and released when complete.

Lumi provides scalable storage and governance for petabytes of operational and subsurface data. Delfi provides on-demand compute for simulation, processing, and model training, with no ceiling on capacity.

The shift from evaluating three development scenarios to evaluating three hundred has a dramatic impact on our customers' understanding of development risk and is only possible when compute is no longer a constraint. Our platforms remove that constraint.

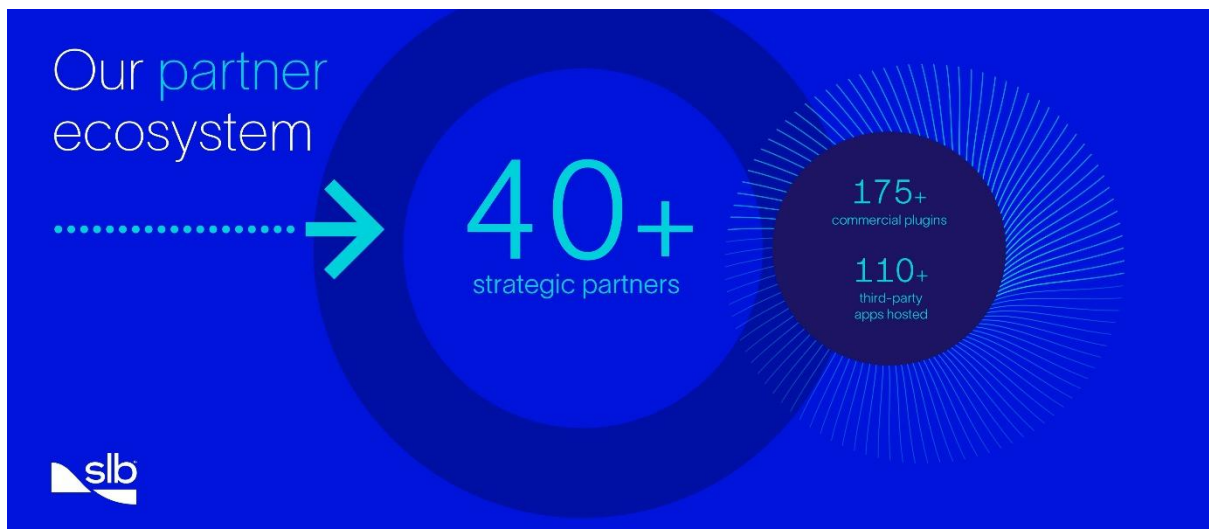
Rakesh has described Tela and the domain foundation models within Lumi, but it is worth stating plainly what this means.

Every agentic AI deployment in this industry depends on the quality of the models, the quality of the data on which they are trained, and the quality of the domain science that governs their outputs. We have all three.

Our models are trained on everything we know and allow customers to enrich them with all that they know. They are grounded in physics not just pattern recognition. And they are deployed within an agentic framework that can act, not just advise.

When we describe our digital platforms, we are describing something quite specific. Not a collection of point solutions. Not desktop applications moved to the cloud. An integrated, open, secure, and scalable platform environment underpinned by the industry's deepest data architecture, powered by domain-native AI, and designed to serve the full lifecycle of an upstream asset.

There is no other platform in the energy industry that offers this combination of workflow depth, data breadth, domain intelligence, extensibility, and enterprise-grade infrastructure. That is the position we have created, and it is the position we will extend.



The final aspect is our partner ecosystem. Let me dive deeper into that.

I should be clear: we did not build all of this by ourselves.

We have more than forty strategic partners contributing capability across the platform that I will talk about in a moment. But the ecosystem extends well beyond strategic partnerships. More than 175 commercial plug-ins are available on our platforms, developed by third parties who have built on our open APIs and frameworks.

And more than 110 third-party applications are hosted on the platform, accessible to our customers alongside our own.

This matters for two reasons.

First, it is evidence that openness is a reality for us, not just a philosophy. Developers and technology companies are investing their own resources to build on our platforms because the customer reach, the data environment, and the commercial opportunity justify that investment. That is the hallmark of a genuine platform ecosystem.

Second, it deepens the moat. Every plug-in, every third-party application, every partner integration increases the value a customer derives from the platform. The ecosystem compounds our own investment and makes the platform more valuable in ways we do not have to build or fund ourselves.



In addition to the other vendors who have brought their IP to our platform, our strategic partner ecosystem above that network is structured in six layers, each serving a distinct function in the platform.

We have foundational partnerships with the hyperscalers: Microsoft, Google, and Amazon Web Services. They are drawn to us as the market leader in the domain, we are drawn to them for their modern cloud infrastructure on which Delfi and Lumi run.

Multi-cloud support is not a convenience but a requirement. Our customers operate all over the world, many with strict data residency constraints. Being cloud-agnostic means we can deploy wherever the customer needs us.

We can also deploy on the edge, using our Agora Edge AI platform. Agora addresses the real-time demands of remote environments where connectivity, latency, cybersecurity and operational continuity directly affect performance.

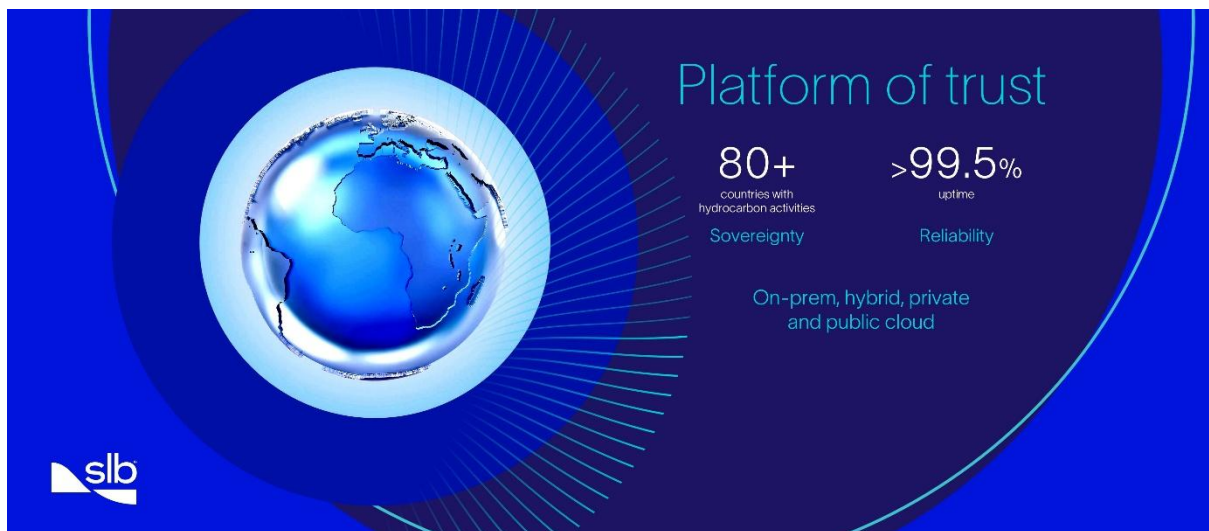
I won't expand upon the other layers. The point of this slide is not the number of logos. We could have added many more. It's the architecture.

Every layer is deliberate. Every partner is best-in-class in their domain.

And together, they create a platform that is comprehensive without being closed.

Let's hear from some of these valued partners.

[SLB partners video plays]



I want to emphasize on something that comes from these videos and that is one of the most significant barriers to entry in this market: trust.


Our customers trust us with their most competitively sensitive data. Seismic surveys that represent hundreds of millions of dollars of acquisition investment. Proprietary reservoir models that underpin multi-billion-dollar development decisions. Real-time operational data from producing assets.

This is data that governments regulate, that boards scrutinize, and that our customers' competitors would love to see. We deliver our trusted platform across more than eighty countries – essentially everywhere that oil and gas is found – each with its own regulatory framework, its own data and technology requirements. And in many cases its own constraints on which cloud infrastructure is permissible or even available.

Solving for these constraints at scale requires multi-cloud capability, on-premises deployment options, and a deep operational understanding of the legal and political landscape in every market we serve. A national oil company in the Middle East has fundamentally different requirements from an independent in the U.S., or a major operating in deepwater Brazil. We serve all of them, whatever their infrastructure constraints.

In addition to sovereignty, reliability reinforces our customers' trust in SLB. We're well over 99.5% uptime. In many cases, we continue delivering the service when the hyperscaler goes down: It is possible precisely because we operate across multiple cloud providers. If one goes down, we move over to another. Our customers' workflows do not stop because a data center in a single region has an outage.

For an operator running real-time production surveillance or time-critical drilling operations, that resilience is a condition of adoption.



If trust is the foundation, scale is the outcome.

A scale we have achieved thanks to more than \$3 billion dollars of research and development spend on digital since 2016 and three hundred and 90 U.S. patents granted in the last five years alone.

The scale of what we have built is worth dwelling on for a moment. These numbers are not projections; they are the current state of the business.

At the center of this slide, you can see more than ninety percent of global production is simulated or modelled using at least one of our digital solutions.

I'll let that sink in for a moment. That is not a market share statistic but a measure of how deeply embedded our technology is in the decisions that Rakesh talked about and that govern the world's hydrocarbon output.

Around it, the operational footprint. More than half a million feet drilled each quarter using SLB automation technology. Over two billion API calls across the platform in 2025; a proxy for the volume of machine-to-machine interaction and the workflows executing continuously across our infrastructure.

Over forty-five million CPU hours consumed in Q1 this year alone, and growing, as customers unlock the simulation, processing, and AI workloads we discussed earlier. Supported by over 2,600 petro-technical experts, the largest team in the industry.

The position is built. The question now is how fast and how far we grow from it.

Now, let me hand back to Rakesh to describe the market opportunity.